

A QUICK-AND-DIRTY CHECK FOR A ONE-DIMENSIONAL ACTIVE SUBSPACE

Paul G. Constantine
Colorado School of Mines, Golden, CO

paul.constantine@mines.edu
<http://inside.mines.edu/~pconstan>

Abstract

Most engineering models contain several parameters, and the map from input parameters to model output can be viewed as a multivariate function. An active subspace is a low-dimensional subspace of the space of inputs that explains the majority of variability in the function. Here we describe a quick check for a dominant one-dimensional active subspace based on computing a linear approximation of the function. The visualization tool presented here is closely related to *regression graphics* [4, 3], though we avoid the statistical interpretation of the model. This document will be part of a larger review paper on active subspace methods.

We can run a quick-and-dirty check to see if a given model has a dominant one-dimensional active subspace—i.e., a direction in the space of model inputs that explains the majority of the variability in the model. After we spell out the steps of this check, we'll describe what it's looking for. Here's what we need to get started:

- *A scalar quantity of interest.* The model should have a single output $f = f(\mathbf{x})$ that depends on the m input parameters denoted $\mathbf{x} = (x_1, \dots, x_m)$. For example, this might be the lift of an airfoil, the maximum power of a battery, or the average temperature over a surface.
- *A range for each input parameter.* Find a lower bound x_i^l and an upper bound x_i^u such that $x_i^l \leq x_i \leq x_i^u$ for each input parameter x_i . This may be a 10% perturbation around some nominal value.
- *Enough time and computing power to run a few model evaluations.* To fit the linear model below, we need to evaluate the model f about three-to-four- m times.

Given these, we run the following procedure.

A quick-and-dirty check for an active subspace

1. Draw N samples $\hat{\mathbf{x}}_j$ uniformly from $[-1, 1]^m$.
2. Let $\mathbf{x}_j = \frac{1}{2} [(\mathbf{x}^u - \mathbf{x}^l) \cdot \hat{\mathbf{x}}_j + (\mathbf{x}^u + \mathbf{x}^l)]$, where \mathbf{x}^l is the vector of lower bounds and \mathbf{x}^u is the vector of upper bounds; the dot operation (\cdot) is component-wise multiplication.
3. Compute $f_j = f(\mathbf{x}_j)$.

4. Compute the coefficients of the linear regression model

$$f(\mathbf{x}) \approx \hat{a}_0 + \hat{a}_1 \hat{x}_1 + \cdots + \hat{a}_m \hat{x}_m$$

with least-squares

$$\hat{\mathbf{a}} = \underset{\mathbf{u}}{\operatorname{argmin}} \|\hat{\mathbf{X}}\mathbf{u} - \mathbf{f}\|_2^2,$$

where

$$\hat{\mathbf{X}} = \begin{bmatrix} 1 & \hat{\mathbf{x}}_1 \\ \vdots & \vdots \\ 1 & \hat{\mathbf{x}}_N \end{bmatrix}, \quad \hat{\mathbf{a}} = \begin{bmatrix} \hat{a}_0 \\ \vdots \\ \hat{a}_m \end{bmatrix}, \quad \mathbf{f} = \begin{bmatrix} f_1 \\ \vdots \\ f_N \end{bmatrix}.$$

5. Let $\mathbf{w} = \mathbf{a}' / \|\mathbf{a}'\|$ where $\mathbf{a}' = [\hat{a}_1, \dots, \hat{a}_m]^T$ is the last m coefficients (i.e., the gradient) of the linear regression approximation.
6. Produce a scatter plot of $\mathbf{w}^T \hat{\mathbf{x}}_j$ versus f_j .

Here is a MATLAB implementation of this procedure, assuming there is a function `mymodel.m` that computes the model output f given inputs and vectors `x1` and `xu` containing the lower and upper bounds of the input parameters.

```
% Produce a scatter plot to check for a
% dominant 1-d active subspace.
N = 4*m;
Xhat = 2*rand(m,N);
X = 0.5*(repmat(xu-x1,1,N).*Xhat + repmat(xu+x1,1,N));
for j=1:N, f(j) = mymodel(X(:,j)); end
ahat = [ones(N,1) Xhat'] \ f';
w = ahat(2:m+1)/norm(ahat(2:m+1));
plot(Xhat*w,f,'o');
```

To demonstrate, we apply this procedure to $f(\mathbf{x}) = \exp(x_1 + x_2)$ defined on the square $[-1, 1]^2$. Figure 1a shows the scatter plot using $N = 20$. Notice the apparent relationship between evaluations $\mathbf{w}^T \mathbf{x}_j$ and f_j . In fact, this particular f varies entirely along a one-dimensional subspace defined by the vector $[1, 1]^T$.

The scatter plot in Figure 1a is equivalent to sampling the three-dimensional surface plot—where the z coordinate is $\exp(x_1 + x_2)$ —and rotating the plot such that the evaluations f_j appear to depend on one variable instead of two. The rotated surface plot along with the f_j is shown in Figure 1b. This interpretation is valid more generally for functions of several variables. The check for an active subspace finds an angle from which to view the data set $\{\mathbf{x}_j, f_j\}$ such that a trend emerges in the f_j (if such a trend exists) as a function of a new variable $y = \mathbf{w}^T \mathbf{x}$.

If, instead of the linear regression procedure, we choose the vector \mathbf{w} as the i th column of the $m \times m$ identity matrix, then we create a scatter plot of f_j versus the i th component of \mathbf{x}_j . Such scatter plots are described in Section 1.2.3 of [6]. However, the vector produced by the regression is special. It is the normalized gradient of the linear regression approximation to $f(\mathbf{x})$. If a global, monotonic trend is present in the f —and if the number of samples N is sufficiently large—then the vector \mathbf{w} reveals the direction in the parameter space of this trend. Additionally, the components of \mathbf{w} measure the relative importance of the original parameters \mathbf{x} .

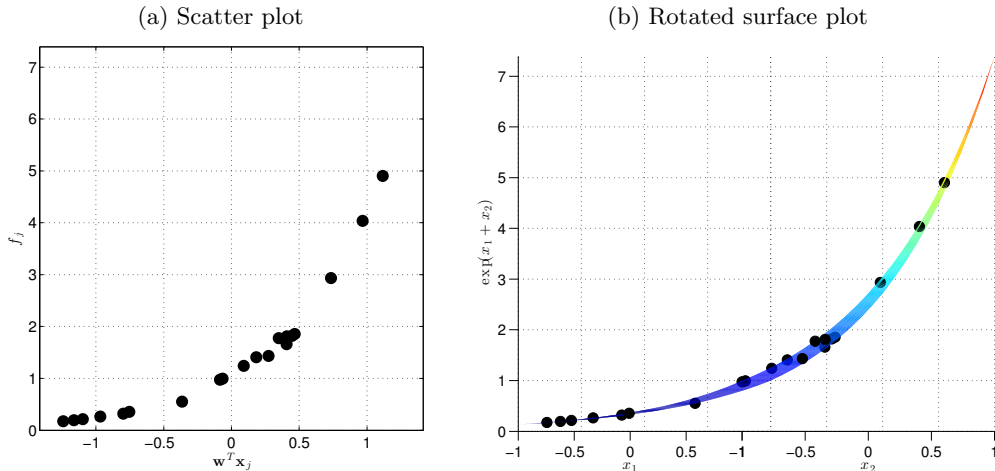


Figure 1: The leftmost figure is the scatter plot produced by the quick check for the active subspace applied to the function $\exp(x_1 + x_2)$. The rightmost figure is a rotated surface plot of the same function.

For the function $\exp(x_1 + x_2)$, we expect the two components of \mathbf{w} to be roughly the same.

This procedure is a subjective diagnostic test based on visualization. The viewer must use the scatter plot to judge if such a trend is present. A trend may be obvious, subtle, or absent. For complex functions, it is unlikely that one can know if such a trend exists without running the check. Fortunately, the check is pretty cheap!

Once a trend is identified, there are several ways to exploit it. For example, if one wants to maximize or minimize $f(\mathbf{x})$, then the vector \mathbf{w} is a direction in the parameter space along which we expect f will increase or decrease. In particular, when the domain of \mathbf{x} is an m -dimensional hypercube, then the vector \mathbf{w} can identify one of the 2^m corners where f is likely to be the largest. If m is large and f is expensive, then this strategy is preferred to evaluating f 2^m times.

If one is attempting to invert f to find parameters that match a given measurement, then this trend can help identify sets of parameters that map to the measurement. If one wishes to predict f at some other values of \mathbf{x} , then this trend can help build a coarse surrogate model on a low-dimensional subspace of the parameter space. The trend may also be used to approximate an integral of f over the parameter space.

Examples.

The check for the active subspace has worked surprisingly well for several complex systems. In particular, it has identified a direction in the parameter space that corresponds to a global monotonic trend in the output. Three representative systems and the corresponding plots are shown in Figures 2, 4, and 3 with a brief description in the captions. More details on these systems can be found in [1], [5], and [2], respectively. We are actively seeking more examples.

There are two common features of these complex models that may explain the perceived trends:

1. When a scientist informally describes the effect of a parameter on a model, it often sounds like the output is monotonic with respect to changes in the parameter; more of x_3 leads to less f . This may suggest a global monotonic trend in the model as a function of the parameter.




Figure 2: We examined a model for heat transfer in a turbine blade given a parameterized model for the heat flux boundary condition representing unknown transition to turbulence. There are 250 parameters characterizing a Karhunen-Loeve model of the heat flux, and the quantity of interest is the average temperature over the trailing edge of the blade. The leftmost figure shows the domain and a representative temperature distribution. The rightmost figure plots 750 samples of the quantity of interest against the projected coordinate. The strong appearance of the linear relationship verifies the quality of the subspace approximation.




Figure 3: We applied the active subspace check to quantify the margins of safety in a multiphysics model of a hypersonic scramjet vehicle. There are seven parameters characterizing the inflow boundary conditions of a reacting, compressible channel flow. The quantity of interest was the exit pressure at the exit of the combustor. The left figure shows a representative compressible flow computation. The right figure shows the scatter plot of 14 samples of the exit pressure.

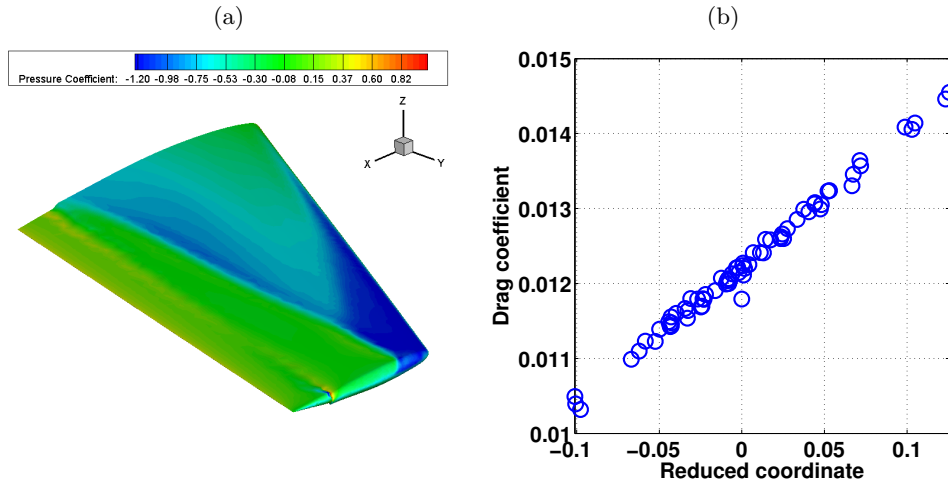


Figure 4: We applied the active subspace check to a design optimization problem in transonic wing design. There were 60 design variables, and the quantity of interest is the drag coefficient computed via computational fluid dynamics. The left figure shows the a representative pressure field for a given design. The right figure shows 70 evaluations of the drag coefficient plotted against the reduced coordinate of the one-dimensional active subspace from the 60-dimensional design space. The nearly linear relationship implies that optimization of the drag coefficient is much easier than initially thought.

2. A model is often given a nominal parameter value corresponding to a theoretical or measured physical case. Back-of-the-envelope estimates of the uncertainty in the input parameters often take the form of a 5-10% perturbation about the nominal value. Despite the complexity of the system, this type of perturbation often results in little change in the model. More importantly, it results in small rate of change. In such cases, the linear model represents the global trends well within the range of the input perturbations.

These explanations are based on experience with several models. Of course, there are many complex systems that do not behave this way. In such cases, the check for the active subspace may lead to nothing useful. However, given the low cost of the check and the potential insights into the model, we think it is well worth the effort.

- [1] P. G. CONSTANTINE, E. DOW, AND Q. WANG, *Active subspace methods in theory and practice: applications to kriging surfaces*, arXiv preprint arXiv:1304.2070, (2013). Version 1.
- [2] P. G. CONSTANTINE, M. EMORY, F. PALACIOS, N. KSEIB, AND G. IACCARINO, *Quantification of margins and uncertainties using an active subspace method for approximating bounds*, in 11th International Conference on Structural Safety & Reliability, 2013.
- [3] R. D. COOK, *Regression graphics: Ideas for studying regressions through graphics*, vol. 482, John Wiley & Sons, 2009.
- [4] R. D. COOK AND S. WEISBERG, *An introduction to regression graphics*, vol. 405, John Wiley & Sons, 2009.

- [5] T. W. LUKACZYK, P. CONSTANTINE, F. PALACIOS, AND J. J. ALONSO, *Active Subspaces for Shape Optimization*, American Institute of Aeronautics and Astronautics, 2014/02/16 2014.
- [6] A. SALTELLI, M. RATTO, T. ANDRES, F. CAMPOLONGO, J. CARIBONI, D. GATELLI, M. SAISANA, AND S. TARANTOLA, *Global sensitivity analysis: the primer*, John Wiley & Sons, 2008.